



US009146613B2

(12) **United States Patent**
Hamdi et al.

(10) **Patent No.:** **US 9,146,613 B2**
(45) **Date of Patent:** **Sep. 29, 2015**

(54) **ADVANCED POE POWER CONTROL**

(56) **References Cited**

(71) Applicant: **Dell Products L.P.**, Round Rock, TX (US)

U.S. PATENT DOCUMENTS

(72) Inventors: **Rabah S. Hamdi**, Spring, TX (US);
Bradley J. Booth, Austin, TX (US)

7,155,622 B2 *	12/2006	Mancey et al.	713/324
7,340,620 B2 *	3/2008	Dove	713/300
7,698,580 B2 *	4/2010	Schindler et al.	713/300
7,849,333 B2 *	12/2010	Schindler	713/300
8,001,397 B2 *	8/2011	Hansalia	713/300
8,132,029 B2 *	3/2012	Alms et al.	713/300
8,205,102 B2 *	6/2012	Kabbara et al.	713/300
8,375,232 B2 *	2/2013	Khan et al.	713/300
2010/0031066 A1 *	2/2010	Geiger et al.	713/300
2010/0205466 A1 *	8/2010	Diab et al.	713/310
2011/0217873 A1 *	9/2011	Diab et al.	439/620.01

(73) Assignee: **Dell Products L.P.**, Round Rock, TX (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 347 days.

* cited by examiner

(21) Appl. No.: **13/738,364**

Primary Examiner — Chun Cao

(22) Filed: **Jan. 10, 2013**

(74) *Attorney, Agent, or Firm* — Haynes and Boone, LLP

(65) **Prior Publication Data**

US 2014/0195831 A1 Jul. 10, 2014

(57) **ABSTRACT**

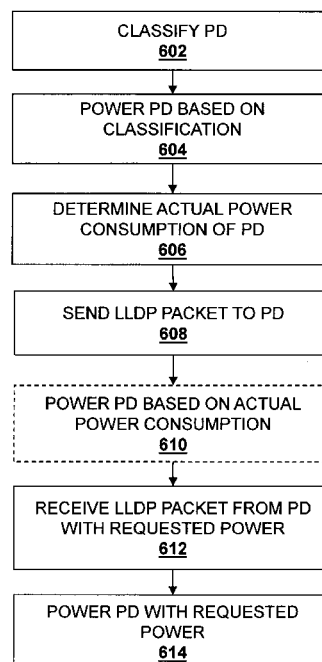
(51) **Int. Cl.**
G06F 1/26 (2006.01)
G06F 1/32 (2006.01)

A power sourcing equipment (PSE) device including a power over Ethernet (PoE) interface. A processor is coupled to the PoE interface. A memory is coupled to the processor and includes instruction that, when executed by the processor, cause the processor to perform a number of functions. The processor may supply power at a first power level to a PD that is coupled to the PoE interface. The processor may then determine an actual power consumption of the PD. The processor may then send a first Link Layer Discovery Protocol (LLDP) packet to the PD over the PoE interface that includes first proposed power level information that is based on the actual power consumption of the PD. Then processor may then supply power to the PD at a second power level that is different from the first power level.

(52) **U.S. Cl.**
CPC **G06F 1/3296** (2013.01); **G06F 1/266** (2013.01); **G06F 1/325** (2013.01); **G06F 1/3215** (2013.01); **Y02B 60/43** (2013.01); **Y02B 60/44** (2013.01)

(58) **Field of Classification Search**
CPC G06F 1/26; G06F 1/266
USPC 713/300
See application file for complete search history.

20 Claims, 9 Drawing Sheets



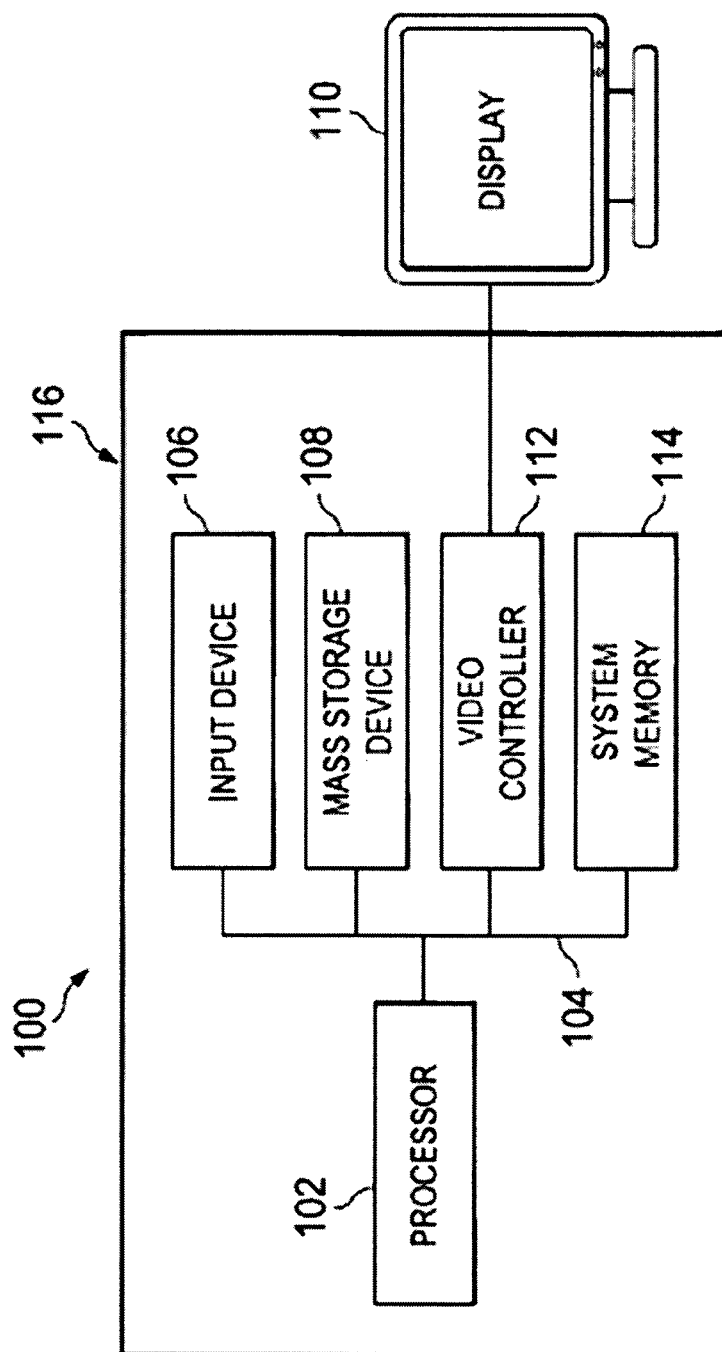


FIG. 1

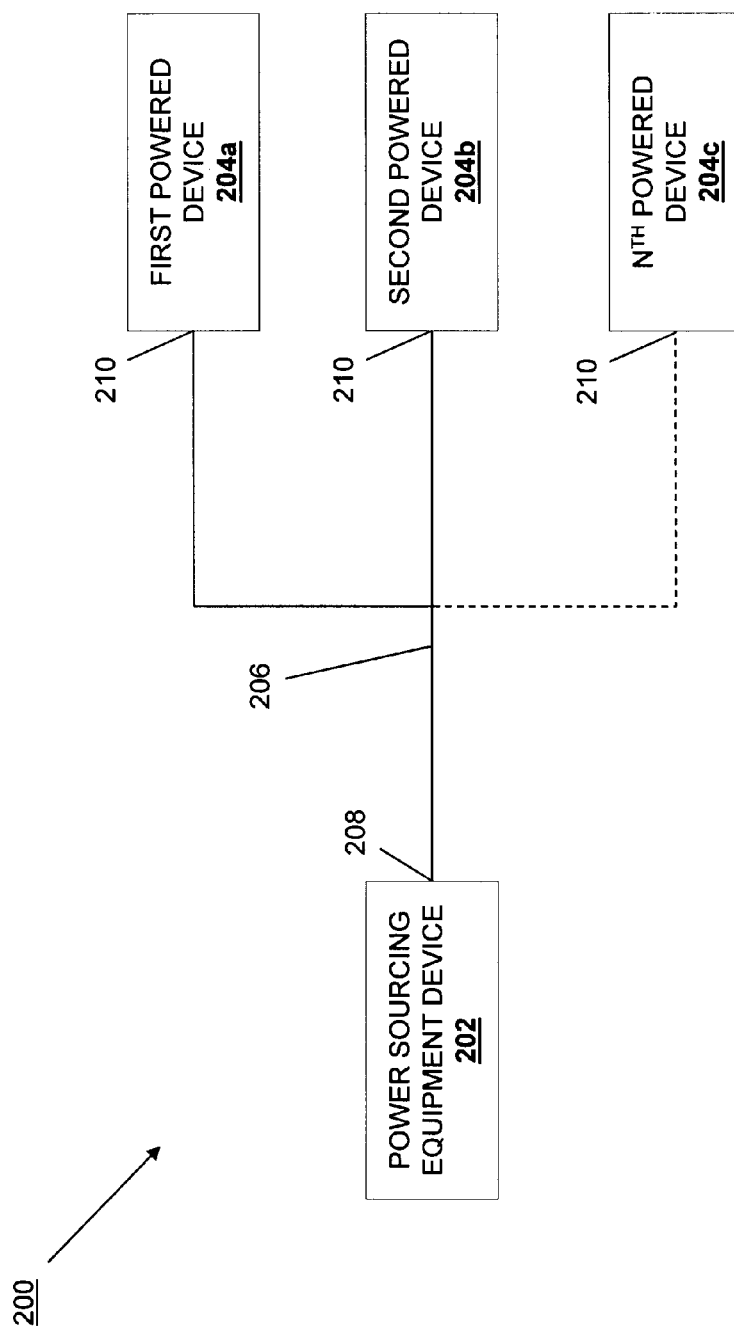


FIG. 2

300




ETHERNET FRAME

Destination MAC Address <u>302</u>	Source MAC Address <u>304</u>	Ethertype <u>306</u>	Chassis ID TLV <u>308</u>	Port ID TLV <u>310</u>	TTL TLV <u>312</u>	Optional TLVs <u>314</u>	End of LLDP PDU TLV <u>316</u>	Frame Check Sequence <u>318</u>
01:80:c2:00:00:00 01:80:c2:00:00:03 01:80:c2:00:00:0e	Station's MAC address	0x88cc	Type = 1	Type = 2	Type = 3	Additional TLVs	Type = 0 Length = 0	FCS

LLDPDU

FIG. 3

400 

Type <u>402</u>	Length <u>404</u>	Organizationally Unique Identifier (OUI) <u>406</u>	Organizationally Defined Subtype <u>408</u>	Organizationally Defined Information String <u>410</u>
7 bits	9 bits	24 bits	8 bits	0-507 octets

FIG. 4

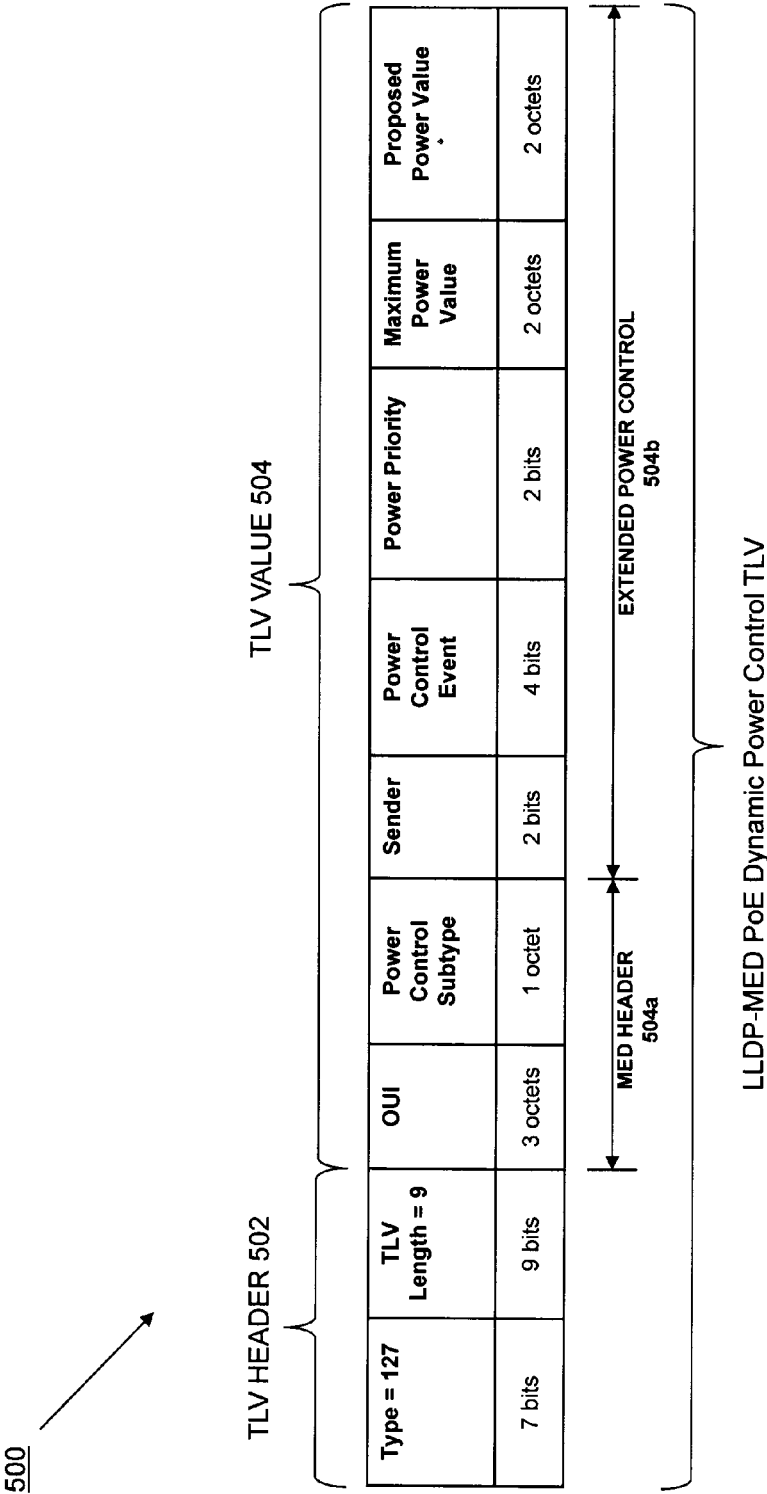


FIG. 5a

Field Name	Values	Description
Type	127	Organizationally specific TLV
Length	9	Length of TLV value
TIA OUI	TBD	Organizationally unique identifier
Power Control Subtype	TBD	Subtype
Sender	PSE = 0 PD = 1	TLV originator
Power Control Event	0 = No event 1 = Proposed Max Power Level 2 = Normal power state 3 = Standby power state 4 = Sleep power state 5 = Remove power 6-15 = TBD	Event or control message
Power priority	0 = Low 1 = High 2 = Mandatory	Power priority and importance level
Maximum power value	0-120 Watts	Maximum allowed power value based on classification, or maximum power value the PD device requires and shall be less than the classifications
Proposed power value	0-120 Watts	Power value conveyed, in 0.1W resolution

FIG. 5b

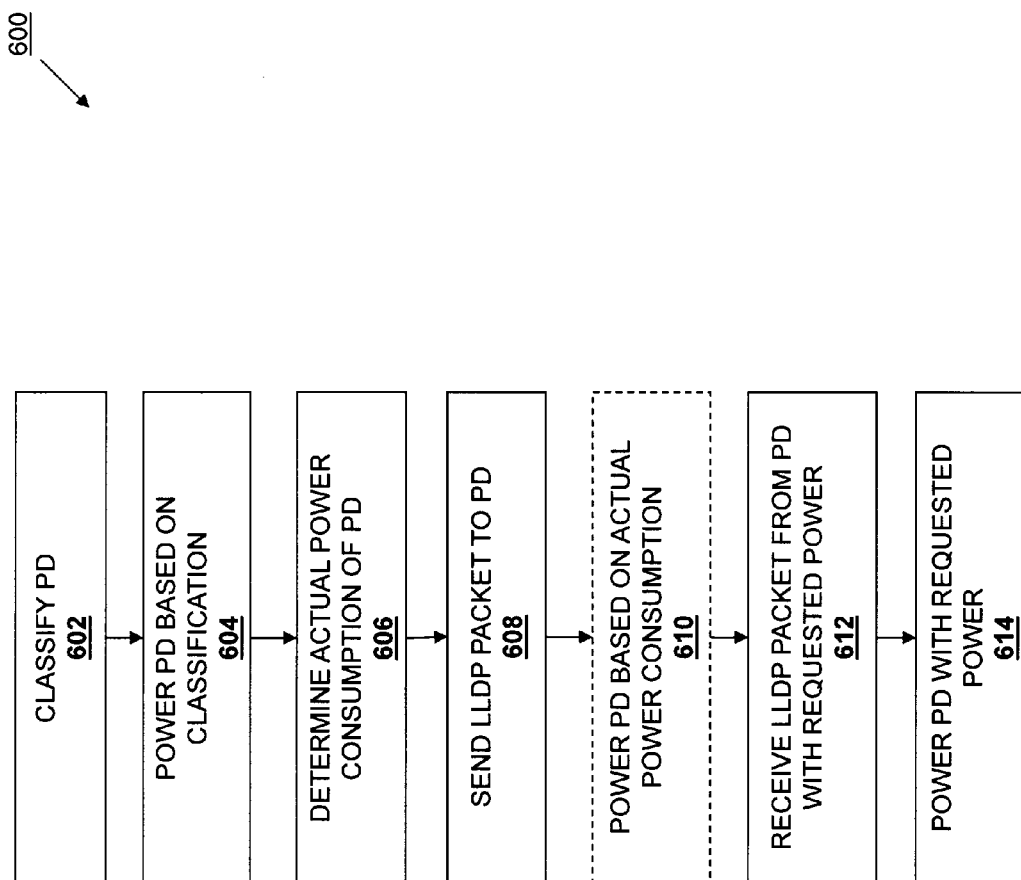


FIG. 6a

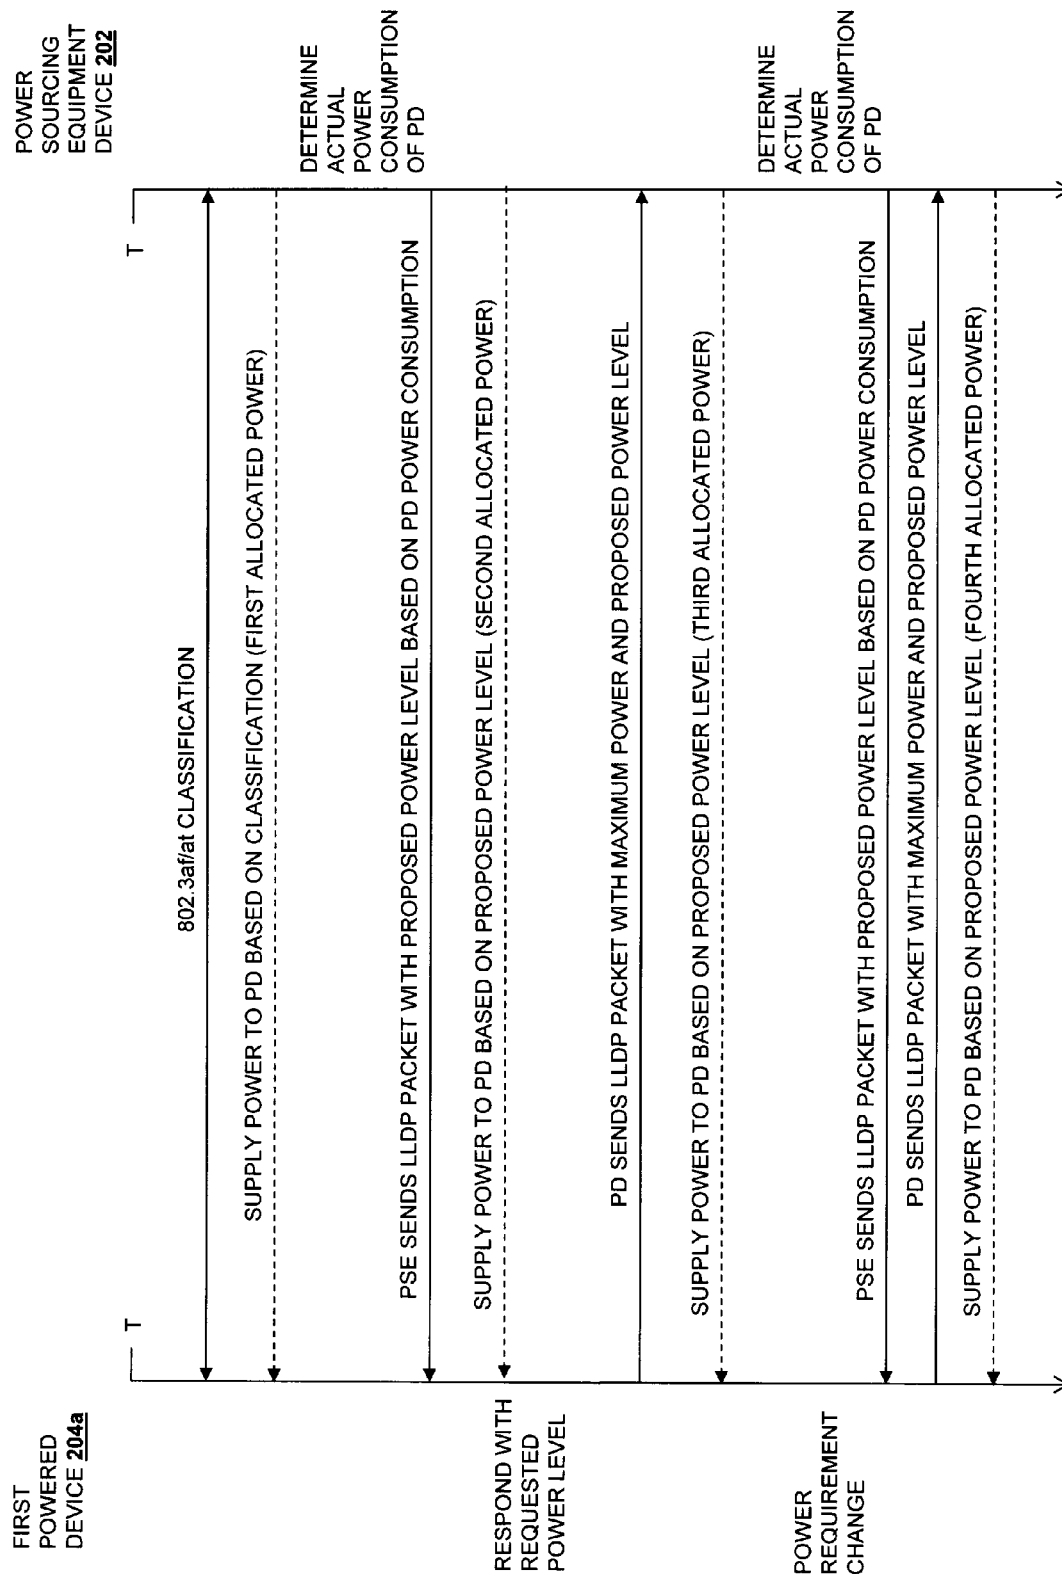


FIG. 6b

700 ↘



FIG. 7

BACKGROUND

The present disclosure relates generally to information handling systems, and more particularly to powering information handling systems using advanced Power over Ethernet (PoE) power control.

As the value and use of information continues to increase, individuals and businesses seek additional ways to process and store information. One option is an information handling system (IHS). An IHS generally processes, compiles, stores, and/or communicates information or data for business, personal, or other purposes. Because technology and information handling needs and requirements may vary between different applications, IHSs may also vary regarding what information is handled, how the information is handled, how much information is processed, stored, or communicated, and how quickly and efficiently the information may be processed, stored, or communicated. The variations in IHSs allow for IHSs to be general or configured for a specific user or specific use such as financial transaction processing, airline reservations, enterprise data storage, or global communications. In addition, IHSs may include a variety of hardware and software components that may be configured to process, store, and communicate information and may include one or more computer systems, data storage systems, and networking systems.

Some IHSs using Power over Ethernet (PoE) technology to send power and data to other IHSs and/or receive power and data from other IHSs. PoE technology provides for the transmission of power safely, along with the data, over Ethernet cabling. The original Institute of Electrical and Electronics Engineers (IEEE) 802.3af standards provide up to 15.4 watts (W) of DC power (minimum 44 volts (V) direct current (DC) and 350 milliamps (mA)), while the updated IEEE 802.3at standards (also known as PoE+) provides up to 25.5 W. The IEEE 802.3af and IEEE 802.3at standards provide for the classification of powered devices (PDs) based on a predefined control protocol or hardware classification, and the power sourcing equipment (PSE) device statically assigns the power level of the power that will be provided to the PD based on the amount designated in the standards for the PD's classification. There are currently 5 classifications: Class 0, the default classification, provides for devices that draw a current of 0-4 mA and a power range of 0.44-12.94 W. Class 1 provides for devices that draw a current of 9-12 mA and a power range of 0.44-3.84 W. Class 2 provides for devices that draw a current of 17-20 mA and a power range of 3.84-6.49 W. Class 3 provides for devices that draw a current of 26-30 mA and a power range of 6.49-12.95 W. Class 4, used by 802.3at devices, provides for devices that draw a current of 36-44 mA and a power range of 12.95-25.5 W.

The static assignment by the PSE device of the power level that will be provided to a PD based on the classification discussed above raises a number of issues. Such static assignments result in the PSE device statically allocating the maximum amount of power needed by a PD of that class when operating (e.g., a PD classified as a class 3 device results in the PSE statically allocating 12.95 W to that PD even if that PD actually draws a maximum of only 8 W power.) Thus, current PoE power systems are inefficient in their allocation of power due to their static assignment of power to PDs based on a classification that can result in the PSE assigning more power to a PD than is necessary.

Accordingly, it would be desirable to provide an improved PoE power control system.

According to one embodiment, a power sourcing equipment (PSE) device includes a power over Ethernet (PoE) interface; a processor coupled to the PoE interface; a memory coupled to the processor and including instruction that, when executed by the processor, cause the processor to: supply power at a first power level to a PD that is coupled to the PoE interface; determine an actual power consumption of the PD; send a first Link Layer Discovery Protocol (LLDP) packet to the PD over the PoE interface that includes first proposed power level information that is based on the actual power consumption of the PD; and supply power to the PD at a second power level that is different from the first power level.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a schematic view illustrating an embodiment of an information handling system.

FIG. 2 is a schematic view illustrating an embodiment of a power over Ethernet (PoE) system.

FIG. 3 is a chart illustrating an embodiment of an Ethernet frame.

FIG. 4 is a chart illustrating an embodiment of an organizationally specific type-length-value (TLV) format.

FIG. 5a is a chart illustrating an embodiment of a link layer discovery protocol (LLDP) media endpoint discovery (MED) PoE power control TLV.

FIG. 5b is a chart illustrating an embodiment of examples of values for the LLDP MED PoE power control TLV of FIG. 5a.

FIG. 6a is a flow chart illustrating an embodiment of a method for providing power.

FIG. 6b is an interaction graph of a PSE device and a PE operating according to the method of FIG. 6a.

FIG. 7 is a schematic view illustrating an embodiment of a PoE system.

DETAILED DESCRIPTION

For purposes of this disclosure, an IHS may include any instrumentality or aggregate of instrumentalities operable to compute, classify, process, transmit, receive, retrieve, originate, switch, store, display, manifest, detect, record, reproduce, handle, or utilize any form of information, intelligence, or data for business, scientific, control, entertainment, or other purposes. For example, an IHS may be a personal computer, a PDA, a consumer electronic device, a display device or monitor, a network server or storage device, a switch router or other network communication device, or any other suitable device and may vary in size, shape, performance, functionality, and price. The IHS may include memory, one or more processing resources such as a central processing unit (CPU) or hardware or software control logic. Additional components of the IHS may include one or more storage devices, one or more communications ports for communicating with external devices as well as various input and output (I/O) devices, such as a keyboard, a mouse, and a video display. The IHS may also include one or more buses operable to transmit communications between the various hardware components.

In one embodiment, IHS 100, FIG. 1, includes a processor 102, which is connected to a bus 104. Bus 104 serves as a connection between processor 102 and other components of IHS 100. An input device 106 is coupled to processor 102 to provide input to processor 102. Examples of input devices may include keyboards, touchscreens, pointing devices such as mice, trackballs, and trackpads, and/or a variety of other

3

input devices known in the art. Programs and data are stored on a mass storage device **108**, which is coupled to processor **102**. Examples of mass storage devices may include hard discs, optical disks, magneto-optical discs, solid-state storage devices, and/or a variety of other mass storage devices known in the art. IHS **100** further includes a display **110**, which is coupled to processor **102** by a video controller **112**. A system memory **114** is coupled to processor **102** to provide the processor with fast storage to facilitate execution of computer programs by processor **102**. Examples of system memory may include random access memory (RAM) devices such as dynamic RAM (DRAM), synchronous DRAM (SDRAM), solid state memory devices, and/or a variety of other memory devices known in the art. In an embodiment, a chassis **116** houses some or all of the components of IHS **100**. It should be understood that other buses and intermediate circuits can be deployed between the components described above and processor **102** to facilitate interconnection between the components and the processor **102**.

Referring now to FIG. 2, an embodiment of a power over Ethernet (PoE) system **200** is illustrated. The PoE system **200** includes a power sourcing equipment (PSE) device **202** coupled to a plurality of powered devices (PDs), such as the first PD **204a**, the second PD **204b**, and up to the NTH PD **204c** of the illustrated embodiment. In an embodiment, any of the PSE device **202**, the first PD **204a**, the second PD **204b**, and up to the NTH PD **204c** may be the IHS **100** and/or include some or all of the IHS components of the IHS **100**, discussed above with reference to FIG. 1. For example, the PSE device **202** may be a switch, a bridge, and/or a variety of other network IHSs known in the art. In another example, the PDs **204a**, **204b**, and **204c** may be network access points (e.g., wireless access points), IP telephony devices, monitoring devices (e.g., camera systems), point of sale devices, and/or a variety of other PDs known in the art. As discussed in more detail below, the PSE device **202** may include a processor and a memory that includes instructions that, when executed by the processor, cause the processor to provide power and data to the PDs **204a**, **204b**, and **204c**, determine the actual power consumption of the PDs **204a**, **204b**, and **204c**, communicate with the PDs **204a**, **204b**, and **204c**, and provide any of the other PSE device functions discussed below. As also discussed below, any of the PDs **204a**, **204b**, and/or **204c** may include a processor and a memory that includes instructions that, when executed by the processor, cause the processor to determine the power requirements of the PD, communicate with the PSE device **202**, and provide any of the other PD functions discussed below.

The PSE device **202** may be coupled to each of the first PD **204a**, the second PD **204b**, and the NTH PD **204c** through one or more Ethernet cables **206** that couple to the PSE device **202** through one or more PoE interfaces **208** and that couple to the PDs **204a**, **204b**, and **204c** through PoE interfaces **210**. While not illustrated, one of skill in the art will recognize that the PSE device **202** may be coupled to a network (e.g., the Internet), a data source (e.g., a server), as well as a power source (e.g., an Alternating Current (AC) power source), and thus may include components for providing data from the network or data source and providing power from the power source safely via the PoE interface **208** and through the Ethernet cable(s) **206**, as discussed in further detail below. Similarly, while not illustrated, one of skill in the art will recognize that the PDs **204a**, **204b**, and **204c** may include components for extracting data and power sent over the Ethernet cable(s) **206** from the PSE device **202** and received through the PoE interfaces **210**, as discussed in further detail below.

4

Referring now to FIG. 3, an embodiment of an Ethernet frame **300** is illustrated that includes a Link Layer Discovery Protocol (LLDP) Data Unit (LLDPDU). The Ethernet frame **300** includes a Destination Media Access Control (MAC) Address **302** that may include a special multicast address (e.g., “01:80:c2:00:00:00”, “01:80:c2:00:00:03”, “01:80:c2:00:00:0e” in the illustrated embodiment) that will prevent an 802.1D compliant bridge from forwarding the Ethernet frame **300**. The Ethernet frame **300** also includes a Source MAC Address **304** that may include the MAC address of the device sending the Ethernet frame **300**. The Ethernet frame **300** also includes an Ethertype **306** that may include the protocol of the payload of the Ethernet frame **300** (e.g., “0x88cc” in the illustrated embodiment.) The Ethernet frame **300** also includes a Chassis identification type-length-value (TLV) **308**, a port identification TLV **310**, a Time To Live (TTL) TLV **312**, one or more Optional TLVs **314**, and an End of LLDP PDU TLV **316**, together which provide the LLDPDU. The Ethernet frame **300** also includes a frame check sequence **318** that may be used for error detection. As discussed below, the Optional TLVs **314** may be used by the PSE device **202** and the PDs **204a**, **204b**, and/or **204c** to convey power requirements and dynamically control power.

Referring now to FIG. 4, an embodiment of an organizationally specific TLV **400** is illustrated that may be used to convey power requirements and dynamically control power. The organizationally specific TLV **400** includes a type **402** provided by seven bits that identify the type of TLV (e.g., the Optional TLVs used in an embodiment of the present disclosure are type **127**). The organizationally specific TLV **400** includes a length **404** provided by nine bits that identify the length of the TLV. The organizationally specific TLV **400** includes an Organizationally Unique Identifier **406** and an Organizationally Defined Subtype **408**, which are provided by twenty-four bits and eight bits, respectively, and are discussed in further detail below with regard to their use in conveying power requirements and dynamically controlling power. The organizationally specific TLV **400** includes a Organizationally Defined Information String **410** that may be provided by 0-507 octets and are discussed in further detail below with regard to their use in conveying power requirements and dynamically controlling power.

Referring now to FIGS. 5a and 5b, an embodiment of an LLDP-MED (Media Endpoint Discovery) PoE Dynamic Power Control TLV **500** is illustrated. The LLDP-MED PoE Dynamic Power Control TLV **500** includes a TLV header **502** including the Type (e.g., “127”) and Length (e.g., “9”) of the TLV. The LLDP-MED PoE Dynamic Power Control TLV **500** includes a TLV Value **504** that provide an MED header **504a** and extended power control information **504b**. The MED header **504a** includes an Organizationally Unique Identifier and a Power Control Subtype that may be defined consistently such that the PSE and PDs recognize the LLDP-MED PoE Dynamic Power Control TLV **500**. In an embodiment, the Power Control Subtype may provide for different power control methods that may vary based on the power control environment. For example, a Power Control Subtype may provide a first power control method appropriate for a data center, a second power control method appropriate for an automobile, a third power control method appropriate for when an intermediate device is present between the PSE device **202** and a PD (discussed in further detail below), and/or other power control methods for a variety of other power control environments known in the art.

The extended power control information **504b** may include variety of information that is detailed further in FIG. 5b. In the illustrated example, the extended power control information

5

504b includes sender information that details the originator of the LLDP-MED PoE Dynamic Power Control TLV **500** (e.g., a “0” for a PSE device and a “1” for a PD). The extended power control information **504b** also includes power control event information that provides an event or control message (e.g., “0” for no event, “1” for a proposed maximum power level, “2” for a normal power state, “3” for a standby power state, “4” for a sleep power state, and “5” for removed power). The extended power control information **504b** also includes a power priority (e.g., “0” for low priority, “1” for high priority, and “2” for mandatory priority). The extended power control information **504b** also includes a maximum power value (e.g., 0-120 watts). The extended power control information **504b** also includes a proposed power value (e.g., 0-120 watts).

Referring now to FIGS. **2**, **6a**, and **6b**, an embodiment of a method **600** for providing power is illustrated. The example of the method **600** discussed below references the PSE device **202** and the first PD **204a** of FIG. **2** to illustrate the communications and power control between the PSE device **202** and the first PD **204a**. However, the method **600** may be performed between any or all of the PDs **204a**, **204b**, and **204c** and the PSE device **202** (and between multiple PDs and a PSE device) while remaining within the scope of the present disclosure. The method **600** begins at block **602** where the PD is classified. In an embodiment, the first PD **204a** is connected to the PSE **202** using an Ethernet cable **206** connected to the PoE interface **208** on the PSE device **202** and the PoE interface **210** on the first PD **204a** and, in response, the first PD **204a** is classified by the PSE device **202** according to 802.3af or 802.3at standards. For example, the PSE device **202** may provide a pulse to measure a resistance in the first PD **204a** that indicates which class (e.g., classes 0-4) the first PD **204a** belongs to, and that class is used to determine what static power level should be allocated to the first PD **204a**.

The method **600** then proceeds to block **604** where the PD is powered based on the classification. In an embodiment, the PSE device **202** is coupled to a power supply that provides the PSE device **202** a power budget that the PSE device **202** may allocate to one or more of the PDs **204a**, **204b**, and **204c**. For example, the PSE device **202** may be operable to calculate the power budget from the total power supplied to the PSE device **202** and available for the PDs **204a**, **204b**, and **204c**. At block **604**, the PSE device **202** supplies the static power level according to the classification performed in block **602**, thus reducing the power budget available for other PDs. As illustrated in FIG. **6b**, the supplying of power to the PD based on the classification provides a first allocated power to the PD. One of skill in the art will recognize blocks **602** and **604** of the method **600** as a prior art classification and PoE power provision. However, the systems and methods of the present disclosure provide for dynamic allocation of power to the PD based on its actual power consumption, power events in the PD, and increased powered demand by the PD, as detailed below.

The method **600** then proceeds to block **606** where the actual power consumption of the PD is determined. In an embodiment, at block **606**, the PSE device **202** monitors the power consumption of the first PD **204a** to determine an actual power consumption of the first PD **204a**. For example, the PSE device **202** may monitor or otherwise determine the average power consumption of the first PD **204a**, the maximum power consumption of the first PD **204a**, and/or a variety of other power consumption details of the first PD **204a**. As discussed above, the first PD **204a** may be classified and then allocated power from the PSE device power budget based on that classification, which can result in more power

6

being allocated to the first PD **204a** than the first PD **204a** is actually consuming (i.e., the first PD **204a** may be consuming only a fraction of the power allocated by the PSE device **202** according to the classification.) Thus, at block **606**, the PSE device **202** may determine that the first PD **204a** is actually consuming a power amount that is less than the first allocated power/static power level provided to the first PD **204a** based on the classification at block **602**.

The method **600** then proceeds to block **608** where an LLDP packet is sent to the PD. In an embodiment, in response to determining that the actual power consumption of the first PD **204a** is different from the power allocated to the first PD **204a** in response to the classification at block **602**, the PSE device **202** may send a LLDP-MED PoE Dynamic Power Control TLV **500** as part of an Ethernet frame **300** to the first PD **204a** in order to propose a new power level that is based on the actual consumption of the first PD **204a**. For example, the LLDP-MED PoE Dynamic Power Control TLV **500** sent in the Ethernet frame **300** by the PSE device **202** may include extended power control information **504b** that includes that the PSE is the sender, a proposed maximum power level as the power control event, a power priority based on the power status of the PSE device **202** (e.g., low priority if the PSE device **202** has surplus power, a high power priority if the PSE device **202** is running out of surplus power, and a mandatory power priority if the PSE device is about to shut down due to lack of power), a maximum power value based on the classification of the first PD **204a**, and a proposed power value based on the actual power consumption of the PD determined at block **606**. The Ethernet frame **300** including the LLDP-MED PoE Dynamic Power Control TLV **500** is sent over the Ethernet cable **206** to the first PD **204a** at block **608**.

The method **600** may then proceed to optional block **610** where the PD is powered based on the actual power consumption determined in block **606**. In an embodiment, optional block **610** may be performed following block **606** and before, concurrently with, or immediately after block **608**. Thus, in some embodiments, the PSE device **202** may immediately begin powering the first PD **204a** based on its actual power consumption determined in block **606**, and the LLDP packet sent to the first PD **204a** may be provided to inform the first PD **204a** that its power allocation from the PSE device **202** has changed from that provided according to the classification at block **602**. Thus, as illustrated in FIG. **6b**, at optional block **610** the PSE device may provide a second allocated power to the first PD **204a** that is based on the actual power consumption of the first PD **204a** and that is less than the first allocated power provided in response to the classification of the first PD **204a**. Furthermore, by sending the LLDP packet to the first PD **204a**, the first PD **204a** is put on notice of the new power allocation (and the limits of that new power allocation) from the PSE device **202**, and in situations where optional block **610** is performed, must then negotiate increased power allocations (or further contribute to surplus power by requesting reduced power allocations) as discussed below. By powering the first PD **204a** according to its actual power consumption, the excess, unused power previously allocated to the first PD **204a** is returned to the power budget of the PSE device **202**. In some embodiments, the PSE device **202** may recalculate the power budget at optional block **610**.

The method **600** then proceeds to block **612** where an LLDP packet with a requested power is received from the PD. In an embodiment, the first PD **204a** may send a LLDP-MED PoE Dynamic Power Control TLV **500** as part of an Ethernet frame **300** to the PSE device **202** in order to request a new power level allocation. In one embodiment, the PSE device **202** may be providing the first PD **204a** power based on the

7

actual power consumption of the first PD **204a** determined in block **606** (and after providing the first PD **204a** the LLDP packet informing the first PD **204a** of the change in power level), and the first PD **204a** may determine that it requires additional power. In another embodiment, the PSE device **202** may have provided the first PD **204a** the LLDP packet indicating to the first PD **204a** of a proposed change in power level allocation, and the first PD **204a** may determine that it requires a different power level allocation than that proposed power level allocation.

For example, the LLDP-MED PoE Dynamic Power Control TLV **500** sent in the Ethernet frame **300** by the first PD **204a** may include extended power control information **504b** that includes that the first PD **204a** is the sender, a power state (e.g., normal, standby, sleep, etc.) as the power control event, a power priority based on the power status of the first PD **204a** (e.g., low priority if the first PD **204a** will operate normally without the requested power, a high power priority if the first PD **204a** is not operate normally without the requested power, and a mandatory power priority if the first PD **204a** is about to shut down due to lack of power), a maximum power value that may include the maximum amount of power the first PD **204a** will need, and a proposed power value that the first PD **204a** is requesting. The Ethernet frame **300** including the LLDP-MED PoE Dynamic Power Control TLV **500** is sent over the Ethernet cable **206** and received by the PSE device **202** at block **612**. In some embodiments, the first PD **204a** may be limited to requesting only as much power as is permitted according to its 802.3af/at classification. In some embodiments, the first PD **204a** may be required to reduce its power consumption to a proposed power level from the PSE device **202** when the power priority in the LLDP-MED PoE Dynamic Power Control TLV **500** is set to mandatory, and must comply with proposed power values in the LLDP-MED PoE Dynamic Power Control TLV **500** received from the PSE device **202**.

The method **600** then proceeds to block **614** where the PD is powered with the requested power. As illustrated in FIG. **6b**, at block **614**, the PSE device may provide a third allocated power to the first PD **204a** that is based on the requested power level/power level proposed by the first PD **204a**. In an embodiment, power states of the first PD **204a** may be defined in the PSE device **202** and associated with power level allocations, and the power state of the first PD **204a** that is provided in the power control event in the extended power control information in the LLDP-MED PoE Dynamic Power Control TLV **500** sent in the Ethernet frame **300** may be used to retrieve the defined power states of the first PD **204a** and look up their associated power requirements/power level allocations. Those associated power requirements may then be used to supply the appropriate amount of power to the PD at block **614**.

As further illustrated in FIG. **6b**, the power requirements of the first PD **204a** may then change, and any or all of the steps described above may be repeated. For example, in the illustrated embodiment, the PSE device **202** may continue to determine the actual power consumption of the first PD **204a**, sending an LLDP packet with a proposed power level along with supplying power to the PD at that proposed power level (a fourth allocated power level), and/or the PD may send an LLDP packet back with another proposed power level based on its new power requirements, all substantially as described above. In some embodiments, after any change in the power supplied to the first PD **204a**, the PSE device **202** may recalculate the power budget.

Thus, a system and method have been described that provide for advanced PoE power control by a PSE device that

8

continually monitors power consumption of a PD to determine the appropriate amount of power to allocate to that PD, along with providing continued negotiation between the PD and the PSE device to adjust the power allocation based on the current power states of the PSE device and the PD. The systems and methods of the present disclosure may be performed by a PSE with any number of PDs in order to efficiently allocate a power available to the PSE device to supply to the PDs. While the details of allocating power discussed above have been discussed based on the actual power consumption determined by the PSE device and changing power needs determined by the PD, the PSE device may be operable to allocate power to PDs based on PSE device port/interface profiles, priorities, events, device classifications, predetermined power states, predetermined power profiles, and/or a variety of other policies/configurations known in the art that may be set by a network administrator.

Referring now to FIG. **7**, an embodiment of a PoE system **700** is illustrated that includes a PSE device **702** coupled to a first PD **704** through an intermediate PoE midspan device **706**. In an embodiment, the operations to supply power from the PoE midspan device **706** to the first PD **704** occur substantially as discussed above according to the method **600**, but with the provision that the PoE midspan device **706** is operable to convey the LLDP packets between the PSE device **702** and the first PD **704**. In some embodiment, the PoE midspan device **706** may supply power to the first PD **704** while the PSE device **702** controls the power supplied to the first PD **704** by the PoE midspan device **706**. In such an embodiment, the combined PSE device **702** and PoE midspan device **706** may operate substantially the same as the PSE device **202** discussed above, with the PoE midspan device **706** operable to determine the actual power consumption of the first PD **704** and provide that information to the PSE device **702**, while conveying the LLDP packets between the PSE device **702** and the first PD **704**.

Although illustrative embodiments have been shown and described, a wide range of modification, change and substitution is contemplated in the foregoing disclosure and in some instances, some features of the embodiments may be employed without a corresponding use of other features. Accordingly, it is appropriate that the appended claims be construed broadly and in a manner consistent with the scope of the embodiments disclosed herein.

What is claimed is:

1. A power sourcing equipment (PSE) device, comprising:
 - a power over Ethernet (PoE) interface;
 - a processor coupled to the PoE interface;
 - a memory coupled to the processor and including instruction that, when executed by the processor, cause the processor to:
 - supply power at an initial power level to a PD that is coupled to the PoE interface;
 - determine a first actual power consumption of the PD;
 - send a first Link Layer Discovery Protocol (LLDP) packet to the PD over the PoE interface that includes first proposed power level information that is based on the first actual power consumption of the PD;
 - supply power to the PD at a first power level that is associated with the first proposed power level information and that is different from the initial power level;
 - receive a second LLDP packet from the PD through the PoE interface subsequent to sending the first LLDP packet, wherein the second LLDP packet includes

9

second proposed power level information that is associated with a second power level that is different from the first power level; and

supply power to the PD at the second power level in response to receiving the second LLDP packet.

2. The PSE device of claim 1, wherein the memory includes instruction that, when executed by the processor, cause the processor to:

classify the PD, wherein the initial power level is supplied to the PD based on the classification of the PD.

3. The PSE device of claim 1, wherein the second LLDP packet includes a power control event Type-Length-Value (TLV) that indicates a power state of the PD, and wherein the power is supplied to the PD at the second power level based on the power state of the PD.

4. The PSE device of claim 3, wherein the power control event TLV includes one of a normal power state, a standby power state, and a sleep power state.

5. The PSE device of claim 1, wherein the second LLDP packet includes maximum power information for the PD.

6. The PSE device of claim 1, wherein the memory includes instruction that, when executed by the processor, cause the processor to:

determine a second actual power consumption of the PD; send a third LLDP packet to the PD over the PoE interface that includes third proposed power level information that is based on the second actual power consumption of the PD; and

supply power to the PD at a third power level that is associated with the third proposed power level information and that is different from the second power level.

7. The PSE device of claim 6, wherein the memory includes instruction that, when executed by the processor, cause the processor to:

receive a fourth LLDP packet from the PD through the PoE interface subsequent to sending the third LLDP packet, wherein the fourth LLDP packet includes fourth proposed power level information that is associated with a fourth power level that is different from the third power level; and

supply power to the PD at the fourth power level in response to receiving the fourth LLDP packet.

8. A power over Ethernet (PoE) system, comprising: a plurality of powered devices (PDs);

a power sourcing equipment (PSE) device coupled to each of the plurality of PDs through a respective Ethernet cable, wherein the PSE device is configured to:

supply power to each of the plurality of PDs at a respective static power level;

determine a first actual power consumption of each of the plurality of PDs;

send a first Link Layer Discovery Protocol (LLDP) packet that includes first proposed power level information to a first PD of the plurality of PDs, wherein the first LLDP packet is sent in response to the first actual power consumption of the first PD being different from the respective static power level of the first PD;

supply power to the first PD at a first power level that is associated with the first proposed power level information and that is different from the respective static power level of the first PD;

receive a second LLDP packet from the first PD subsequent to sending the first LLDP packet, wherein the second LLDP packet includes second proposed

10

power level information that is associated with a second power level that is different from the first power level; and

supply power to the first PD at the second power level in response to receiving the second LLDP packet.

9. The PoE system of claim 8, wherein the PSE device is configured to:

classify each of the plurality of PDs, wherein the respective static power level of each of the plurality of PDs is supplied based on the classification of that PD.

10. The PoE system of claim 8, wherein the second LLDP packet includes a power control event Type-Length-Value (TLV) that indicates a power state of the first PD, and wherein the power is supplied to the PD at the second power level based on the power state of the first PD.

11. The PoE system of claim 10, wherein the power control event TLV includes one of a normal power state, a standby power state, and a sleep power state.

12. The PoE system of claim 8, wherein the second LLDP packet includes maximum power information for the first PD.

13. The PoE system of claim 8, wherein the PSE device is configured to:

determine a second actual power consumption of each of the plurality of PDs;

send a third LLDP packet that includes third proposed power level information to the first PD, wherein the third LLDP packet is sent in response to the second actual power consumption of the first PD being different from the second power level;

supply power to the first PD at a third power level that is associated with the third proposed power level information and that is different from the second level.

14. The PoE system of claim 13, wherein the PSE device is configured to:

receive a fourth LLDP packet from the first PD subsequent to sending the third LLDP packet, wherein the fourth LLDP packet includes fourth proposed power level information that is associated with a fourth proposed power level that is different from the third power level; and

supply power to the PD at the fourth power level in response to receiving the fourth LLDP packet.

15. A method for providing power, comprising: supplying, by a power source equipment (PSE) device, power at an initial power level to a powered device (PD); determining, by the PSE device, a first actual power consumption of the PD;

sending, by the PSE device, a first Link Layer Discovery Protocol (LLDP) packet to the PD that includes first proposed power level information that is based on the first actual power consumption of the PD;

supplying, by the PSE device, power to the PD at a first power level that is associated with the first proposed power level information and that is different from the initial power level;

receiving, by the LLDP packet, wherein the second LLDP packet from the PD subsequent to sending the first LLDP packet, wherein the second LLDP packet includes second proposed power level information that is associated with a second power level that is different from the first power level; and

supplying power, by the PSE device, to the PD at the second power level in response to receiving the second LLDP packet.

11

16. The method of claim **15**, further comprising:
 classifying, by the PSE device, the PD, wherein the initial
 power level is supplied to the PD based on the classification of the PD.

17. The method of claim **15**,
 wherein the second LLDP packet includes a power control
 event Type-Length-Value (TLV) that indicates a power
 state of the PD, and wherein the power is supplied to the
 PD at the second power level based on the power state of
 the PD.

18. The method of claim **15**,
 wherein the second LLDP packet includes maximum
 power information for the PD.

19. The method of claim **15**, further comprising:
 determining, by the PSE device, a second actual power
 consumption of the PD;
 sending, by the PSE device, a third LLDP packet from the
 PSE device to the PD that includes third proposed power

12

level information that is based on the second actual
 power consumption of the PD; and
 supplying, by the PSE device, power to the PD at a third
 power level that is associated with the third proposed
 power level information and that is different from the
 second power level.

20. The method of claim **19**, further comprising:
 receiving, by the PSE device, a fourth LLDP packet from
 the PD subsequent to receiving the third LLDP packet,
 wherein the fourth LLDP packet includes fourth pro-
 posed power level information that is associated with a
 fourth power level that is different from the third power
 level; and
 supplying, by the PSE device, power to the PD at the fourth
 power level in response to receiving the fourth LLDP
 packet.

* * * * *